

AP104TP

APPLICATION

FOR

UNITED STATES LETTERS PATENT

Be it known that I, Aindraís O'Callaghan, a citizen of the United States of America, of 956 Hyde St. #12, San Francisco, CA 94109, have invented new and useful improvements in:

REDEYE REDUCTION OF DIGITAL IMAGES

of which the following is the specification.

RECEIVED
MAR 13 2001
U.S. PATENT & TRADEMARK OFFICE

CERTIFICATION UNDER 37 C.F.R. 1.10

"Express Mail" Mailing Label Number: EL700476173US

Date of Deposit: April 13, 2001

I hereby certify that this patent application is being deposited with the United States Postal Service on this date in an envelope as "Express Mail Post Office to Addressee" service under 37 C.F.R. 1.10 on the date indicated above and is addressed to the Assistant Commissioner for Patents, Washington, DC 20231.



Rosalio Haro

REDEYE REDUCTION OF DIGITAL IMAGES

Inventors: Aindraais O'Callaghan

BACKGROUND OF THE INVENTION

5 Field of the Invention

This invention relates generally to the reduction of redevye in digital images, and specifically to the identification and correction of redevye regions within a digital image.

Description of the Related Art

10 Redeye occurs when a flash positioned near the lens of a camera enters the pupils of a subject's eyes, and then bounces off the eye's blood-rich retinal membranes back to the camera. The resultant image may now have a redevye effect, namely the image of the pupil of the eye appears to be unnaturally red. The redness can be mitigated through optical means or through digital processing.

15 Although most cameras attempt to reduce it, the "redevye" effect appears in many photographs. At the instant the photo is taken, bright light from the flash enters the eye through the pupil and reflects off the retina (the back of the eye) to the camera. The unwanted reflection leaves the subject with unnaturally red, glowing eyes.

20 In most redevye photos, the effect conforms to the following description. The sclera (the white part) is not affected. The iris (colored part) of the eye is not affected. The cornea (the clear part in front of the pupil) is not affected. Although this portion of the eye is clear, a small reflection of the flash from the cornea leaves a natural looking 'glint' in the eye. This glint is superimposed over the pupil. The
25 pupil (the hole in the eye that one sees through) is affected. Instead of appearing dark or black, the pupil is red and luminescent. Regardless of whether the pupil is red or black, the glint in the cornea is visible as a very light area within the pupil.

In digital photos, because the precision of focus and resolution of display are not infinite (and may indeed be very poor), the redevye-affected region may spill over
30 to the regions immediately adjacent to the pupil. This region of blending usually extends one or two pixels out of the pupil, but can be greater, especially if the picture has been blown up to a size greater than its resolution allows for.

The reduction or elimination of redeye can be attempted while the photograph is being taken or after it has been taken. The common approach for redeye reduction adopted in cameras is usually the following: a pre-flash is made before the photo is actually taken, thereby conditioning the eye to lessen the opening size of the pupil diameter. Therefore when the flash of the actual photograph goes off, the reflection of the retina is not so noticeable. Positioning of the flash well above the camera lens can also mitigate the redeye effect. Regardless of whether precaution against redeye was taken by the photographer, redeye may still end up in the resulting picture. Indeed, there are many photos in circulation today that would benefit from the removal of redeye. Therefore, it is valuable to discuss the removal of redeye from a photo after it has been taken.

The technique to remove redeye from a photo involves two phases: redeye pupil location and color replacement. In the redeye pupil location phase, the pupil or pupils exhibiting the redeye effect are found. In the color replacement phase, the color of the pupil or pupils is changed so that they look more like ordinary pupils than redeye pupils.

Pupil location is the more interesting and difficult problem. It can be accomplished manually or automatically. Manual pupil location requires that the user identify the problem region. Automatic pupil location is very difficult to reliably ensure due to possible interference or confusion from other parts of the image. An example of automatic pupil location is shown in U.S. Pat. 6,151,403. A reasonable balance between fully manual and fully automatic pupil location is automatic detection within a manually selected area of interest. This allows the user to crop out any confusing areas of the photo so that the automatic redeye detector can more easily find the elusive pupil or pupils. Examples of redeye reduction are found in the following U.S. patents: 6,134,339; 6,160,923; 6,009,209; 5,596,346; and 5,432,863.

The success of different redeye reduction techniques varies more in the pupil location algorithm than in the color replacement. Poor color replacement may leave the image looking scarcely better than it had before the touch-up. Poor pupil location will surely make the image look worse. Location of the pupil is also more time consuming than color replacement. Speed is crucial to good pupil location.

U.S. Pat. 6,016,354 to Lin et al., incorporated herein by reference, presents the following scheme for redeye reduction. The user selects a small rectangular section of a digital image within which lies at least one redeye pupil. A threshold is applied to the pixels within that section so that pixels with a value above the threshold are considered as candidate redeye pixels and those below are ignored

(the threshold is applied to a measure of redeye likelihood, not just the pixels' R value). A pupil search is conducted on the candidate pixels. This search yields the largest circular collection of pixels with the highest concentration of candidate pixels by least squares. Some extra work addresses the issue of the pupil found containing no more and no less than the entire instance of that redeye pupil. Each pixel in the pupil is replaced with a gray value with 70% of the original's luminescence. This attempts to preserve the pupil's glint while leaving it with a more natural color.

Adobe's Photo Deluxe© Business edition has a redeye reduction function that works quite effectively and has a good balance between speed of execution and effectiveness. While the underlying algorithm details are not known to the Applicant, the following observations were made: (a) it uses the same input as required by Lin et al., (b) the pupil location algorithm seems to be different than the method of Lin et al. in that it appears faster and it attempts to find a second pupil after it finds the first. Its color replacement produces similar results as Lin et al.. The pupil is close to its proper color and the glint appears nearly as bright as it should be.

However, both have drawbacks. Redeye correction is difficult, and nobody does a spotless job. Therefore, the best redeye algorithm is that with the fewest faults. Applicant examined various faults that occur in redeye algorithms. These can be categorized as qualitative and quantitative faults in results, and speed issues. It seems difficult to imagine an algorithm requiring significantly less space for operation since the entire input rectangle must be in memory at once.

The algorithm of Lin et al. required $h*w*\log(\min(h,w))$ time for pupil location where h and w are the sides of the input rectangle. Photo Deluxe© uses about the same amount of time. It achieves an almost instantaneous result with a small input rectangle, but slows down significantly with larger ones and ultimately pops up a warning dialogue box if the input is huge. Because it is likely that the instance of redeye fits within a small rectangle, this behavior isn't of much concern. It would be nice, however, to find an algorithm that is linear in time with respect to n, defined as $h*w$. This algorithm would go over each pixel in the input a constant number of times before deciding where the pupils are.

For the term "quantitative faults" to have any meaning in terms of results, it must be defined. We will call failure to locate a redeye pupil a quantitative error in a redeye correction algorithm. Also, redeye software that mistakenly locates an area as redeye will be said to have a quantitative fault in its result. Because all

techniques will be susceptible to some sort of error in this area, "improvement" may not be a hard and fast term. Lin et al. have two major faults here.

Adobe's Photo Deluxe© not only is able to reduce the redeye effect like the algorithm of Lin et al. but it also attempts to find a second eye. Unfortunately, in cases where only one redeye pupil is in the input rectangle (or when it fails to find an existing second pupil), an unwanted circle is declared a redeye pupil and corrected as such, leaving an unsightly circular blotch of gray on a pimple or a red earring, or other image having round red characteristics.

Qualitative faults in results encompass poor replacement color or strategy, unnecessary obscuring of the pupil's glint, or lack of versatility in high/low brightness/contrast situations. Both Lin et al. and Adobe Photo Deluxe©, seem to do a sub-optimal job of preserving the pupil's glint. Each algorithm comes up with a solid circular area (encompassing this glint) to darken and gray. The shine is preserved at all by the fact that the pupil is not darkened too much. Unfortunately, this can sometimes be noticed when the pupil is not quite dark enough to look real. A better color replacement algorithm is needed. Preferably one that would leave the glint alone and darken only the red pupil area.

Objects of the Invention

Therefore, it is an object of the present invention to provide a color replacement algorithm that does not harm the light glint in the pupil.

It is another object of the present invention to provide a redeye correction system whose execution time is linear in time with respect to the number of pixels in a region to be corrected.

It is another object of the present invention to provide a quick, simplified, and effective redeye correction system.

It is still another object of the present invention to provide a redeye reduction system capable of quickly determining if a second pupil exists within an image region to be corrected.

Summary of the Invention

The present invention uses the semiautomatic approach to reduce the redeye effect using digital image processing techniques. That is, an area of interest is cropped and submitted by a user, and the present invention then automatically locates a first and, possibly, a second pupil within the area.

A color-and-brightness-based parameter of the pixels in an area of interest is first compared with a first threshold value to produce a first matrix of candidate redeye pixels. The matrix is then refined by removing non-cluster regions of candidate pixels. The remaining pixels are then gathered into components of cohesive candidate pixels. As the components are determined, statistical data about each component is generated and stored. The statistical data may include the highest color-and-brightness-based parameter value of the component's pixels, the average color-and-brightness-based parameter value of the component's pixels, the size of the component, and optionally a boundary square population value of the component, which is a measure of the pixel density within the smallest square in which a component fits, the higher the boundary square population value the denser the pixel population.

A second matrix may also be generated by comparing the pixels in the area of interest with a second threshold value lower than the first threshold value.

The component having the pixel with the highest color-and-brightness-based parameter value is identified as a first redeye pupil region. The average color-and-brightness-based parameter value and size parameter value of the remaining components are then checked to see if they lie within respective predefined ranges based on the first redeye pupil's average color-and-brightness-based parameter value and size parameter value. If no component has both parameters within the predefined ranges, then no second pupil is identified. Components that do match the calculated ranges are categorized as candidate components and their respective boundary square population values are compared. The candidate component having the densest population of pixels is designated as a second redeye pupil region.

The identified redeye pupil regions are then dilated beyond their boundary and all pixels within the dilated boundary that meet the above described second threshold value are categorized as candidate redeye pixels. Thus, the boundaries of the redeye regions are redefined. The redefined redeye regions are colored with a dark gray that corresponds to the pixels' original values, and the pixels adjacent the redefined redeye regions are colored with a lighter gray.

Other objects and attainments together with a fuller understanding of the invention will become apparent and appreciated by referring to the following description and claims taken in conjunction with the accompanying drawings.

Brief Description of the Drawings

In the drawings wherein like reference symbols refer to like parts,

Fig. 1 is a diagrammatic illustration of the present invention.

Fig. 2 is an example of a digital image having redeye regions selected for input to the present invention.

5 Figs. 3-7 illustrate various stages of matrix manipulation in accord with the present invention.

Fig. 8 shows the image of Fig. 2 with redeye regions corrected and glint regions preserved.

Description of the Preferred Embodiments

10 The prior art discussed above has a few drawbacks. First, the prior art's selection of the largest, most dense collection of redeye candidate pixels as the redeye region is not always the best way to find the pupil. Applicant has found that other factors, such as average brightness or brightest element may help to distinguish between a bright redeye pupil and a dull red section of a sweater. Second, in finding and eliminating redeye one eye at a time, the prior art throws away valuable information that may be useful in locating additional unwanted redeye regions.

15 The redeye effect usually comes in pairs. One eye, however, is usually looking slightly more directly at the camera's flash than the other due to the fact that the eyes are a couple of inches apart, but pointing the same direction. What this means to the digital photo is that, especially in close-up pictures where the parallax is very great, one eye will be brighter than another. In fact, the less extreme of the two red eyes may be very slight, but still noticeable. In practice, the more severe eye will be very easy to locate. The less severe eye may be very difficult to locate on its own, but can be much easier with information about the first eye. Its size, location, shape, and intensity can all contribute in finding the hard-to-see redeye. Correcting the first eye alone will not only lose this information, but will also contribute to the background red-noise, actually making finding the second eye more difficult since its intensity is likely to be lower and thus more easily lost among other red regions of an image.

20 The present invention combines features of existing redeye correction techniques with additional innovations to produce a redeye reduction algorithm with better results. The input to the algorithm is a rectangular area of interest within a digital image selected by a user. The output is that section corrected for redeye. Furthermore, the running time of the present redeye correction technique is linear with respect to the input size in pixels.

With reference to Fig. 1, the first step 11 is to generate a matrix highlighting a specific characterization of the pixels within the rectangular area of interest. The matrix typically compares each pixel with a color-based parameter threshold, such as looking for a particular red color and brightness intensity, and identifies all the pixels that meet the threshold value. Many parameter matrix generating routines and criteria are known in the art, and of particular interest to the present invention is one based on the chrominance in the red color plane of an altered form of the RGB representation of the image.

In the presently preferred embodiment, matrix generator block 11 first raises each pixel's *R*, *G*, and *B* values (i.e. read, green, and blue values) to the $1/3$ power. This seemingly bizarre transformation of the image has been experimentally determined to make redeye easier to distinguish. The altered RGB data is converted to YCrCb values where *Y* represents luminance information and *Cr* and *Cb* represent chrominance information. A first threshold *T* is assigned the value

$$T = Cr_{avg} + 0.2*(Cr_{max} - Cr_{min})$$

where Cr_{avg} is the average chrominance *Cr* value the rectangular area, Cr_{max} and Cr_{min} are respectively the maximum and minimum chrominance *Cr* values. Matrix generator block 11 generates a matrix by placing a 0 in every corresponding pixel location whose modified *Cr* value is below *T* and placing a 1 for every place above *T*. Each 1 represents a candidate redeye pixel.

The resultant matrix is then preferably applied to a candidate-erosion block 13 to refine the matrix. Candidate-erosion block 13 preferably includes a first operation 13a, in which it removes from consideration any candidate pixel without 3 or more candidates pixels as neighbors, followed by a second operation 13b, in which it remove any remaining candidate pixels that now have fewer than 2 candidates pixels as neighbors. This erosion of the candidate redeye areas favors solid redeye sections. These solid groups of candidate pixels are more likely to be the pupils than similarly sized broken up, or elongated, sections. Also, tiny isolated spots are removed and therefore do not interfere with later computations. These two passes 13a and 13b through the matrix are accomplished in linear time and can be implemented as a variation on cellular automata (considering only pixels in the mask for possible deletion rather than considering other pixels for possible genesis).

Next, a component-connecting block 15 implements a connected component algorithm that identifies each cohesive group of candidate redeye pixels for individual consideration. In essence, component-connecting block 15 discerns components made up of cohesive candidate pixels and stores them in a list.

Preferably, statistical data about each cohesive group of candidate pixels is generated as each component is identified and is stored in the list with components. The statistical data recorded for each component include, for example, the value of the brightest pixel (preferably in terms of the above described color-based parameter), average brightness or average color-based parameter value, and the component size, which may be a cumulative pixel size. This is accomplished in linear time with respect to the number of pixels by considering each pixel once for each statistic sought due to the constant number of statistics being recorded and their simple nature.

First redeye identifier block 17 now targets the component that is most likely to be a redeye pupil. This is determined by picking the component containing the pixel with the highest modified Cr value (i.e. the highest color-based parameter value, which includes a measure of brightness). Applicant has found that this selection metric repeatedly yielded the best results in experiments and was quite uncannily good and linear with respect to the number of components (obviously less than n , where n is the total number of pixels in the entered image region of interest).

Having information about the first pupil in hand makes finding a second pupil easier. Therefore, second redeye identifier block 19 seeks a second pupil among the remaining components. First, a candidate component generator block 19a cycles through the statistical data gathered in block 15 and identifies all remaining components whose size parameter is within a preferred range of 0.5 to 1.5 times the size parameter of the component identified as the first redeye region in block 17, and whose average brightness parameter (or average color-based parameter) is within a preferred range of 0.5 to 1.5 times the corresponding parameter in the first redeye component identified in block 17. The components that meet this criteria form a list of candidate components (drawn from elements of the original component list) whose size and brightness are within certain range constraints determined by the first redeye region. In this way, the second pupil selection is constrained by the characteristics of the first pupil region. It is to be understood that these range constraints may be adjusted for particular applications, and that a tighter range will reduce the number of resultant candidate components while a looser range will increase the number of resultant candidate components.

A determination is made at stage 19b as to whether a second pupil region exists based on the availability of candidate components. If no candidate components are identified in block 19a, i.e. none of the remaining components have

characteristics similar to the first pupil region, then the operation of the second redeye identifier block 19 ends.

Otherwise, target component selector block 19c selects a second pupil region from the identified list of candidate components. This is done by calculating a boundary square population value for each candidate component, which is a measure of the pixel density within the smallest square in which a component fits. Alternatively, the boundary square population value of each component may be calculated, along with the other statistical data, in block 15 as each component of cohesive pixels is defined. Target component selector block 19c identifies the candidate component with the highest boundary square population value as the second pupil region. In this manner, this second pupil is determined to be the candidate component whose bounding square is most densely populated with redeye pixel candidates. These constraints require that the second pupil be close in size to the first pupil and nearly as bright. Picking the most densely packed candidate component favors circular components over oblong and scattered ones in order to pick the most pupil-like candidate. As stated above, if no suitable candidate is found, no second pupil is picked. The search for the second pupil takes linear time according to the number of candidate components found because the data used in calculation has already been computed in linear time and is conveniently stored in the candidate list.

Now that the pupils to be corrected have been found, one more step can refine the data so that the true extent of the pupils might be more closely approximated. It may happen that the redeye pupils contain some very red spots and some slightly red spots. The following processing sorts through pixels in the near vicinity to determine whether they should be corrected. Component boundary adjustment block 21 adjusts the areas of the identified one or two redeye pupils, i.e. target components, that are to be corrected by dilating their area out twice. This may be done by running two iterations of a cellular automaton that marks any pixel whose self or immediate neighbor is marked, but other methods of dilating known in the art are likewise suitable to the present invention.

Pupil redefinition block 23 then determines if the newly encompassed pixels within the dilated components meet a relaxed second threshold value T_2 based on the first threshold value. This may be done several ways. In one approach, original image data is fed through the relaxed second threshold T_2 , which is preferably defined as

$$T_2 = 0.9994 * T.$$

A second matrix of candidate pixels based on T_2 is thus constructed, but could alternatively have been constructed along with the first candidate matrix constructed by matrix generator 11 for time efficiency. In still another embodiment, the second matrix is constructed by examining only the pixels within the dilated components, and all pixels outside the dilated components are designated non-candidate pixels without examination.

Final redeye matrix generator block 25 then creates a final redeye replacement matrix by executing a bit-wise *AND* of the dilated original first matrix generated by block 11 and the new relaxed candidate matrix generated by block 23. This allows slightly less red neighbors to be included in the final replacement. If they had been included from the start in the formation of the first matrix, finding the pupils in the first place would have been more difficult because of the extra noise that would have resulted from the use of the lower threshold value.

Now that the actual pupils have been discovered, image modifier block 27 corrects the redeye pixels. Though the partially red pixels have been picked up by blocks 21-25, there may be a one-pixel layer of half-redeye pixels. These are pixels whose borders encompassed some redeye and some non-redeye in continuous space and who have been blurred by the discretizing operation of making the photograph digital. These pixels need to have their red removed but would look bad if they were darkened too much. This is where prior art redeye correction algorithms decide their darkness of pupil vs. lightness of glint compromise. Instead, the present invention replaces the image pixels corresponding to the final redeye replacement matrix generated in block 15 with a dark gray color while image pixels bordering it are replaced with a light gray color. In this way, the glint is a bright gray spot in the center of the eye, while the neighboring pixels of the pupil can be very dark. No pixel, however, has any of its original red color. Pixels between the pupil and iris form a feathered border to the corrected region so that no noticeable hard edges between the corrected and uncorrected sections are produced as artifacts. In this way, the replacement pixels are given gray values equal to 0.6 and 0.9 times their original brightness (measured as the sum of the R, G, and B values of the pixel in the original digital image), respectively.

The present invention may be implanted in hardware, software, or a combination of both.

An example of the present redeye reduction technique applied to a digital image is shown in Figs. 2-8. With reference to Fig. 2, a sample digital image shows first 33 and second 35 redeye regions within a user defined input region of interest represented by a selection rectangle 37. In close proximity to redeye

regions 33 and 35 are respective glint regions 39 and 38. Also within selection rectangle 37, is part of a drinking glass 41 and part of a cloth 43. Drinking glass 41 has a reddish tone and cloth 43 is bright red in color such that both glass 41 and cloth 43 constitute sources of confusion in attempting to identify redeye regions 39 and 38.

The partial image defined by selection rectangle 37 is applied to matrix generator block 11 of Fig. 1, which as explained above, produces a matrix 50 of candidate and non-candidate pixels shown in Fig. 3. Candidate pixels are shown as white regions in Fig. 3, and non-candidate pixels are defined by darkening cross-hatch lines. Candidate pixels 53 correspond to redeye region 33. Candidate pixels 55 correspond to redeye region 35. Candidate pixels 51 correspond to reddish tones in glass 41, and candidate pixels 57 correspond to bright red cloth 43.

Matrix 50 is then applied to candidate-erosion block 13, which removes non-cluster regions of candidate pixels. The resultant, refined matrix 60 is shown in Fig. 4. All elements similar to those of Fig. 3 have similar reference characters and are defined above. As explained above, candidate-erosion block 13 first remove all candidate pixels having less than 3 neighboring candidate pixels, and then removes all remaining candidate pixels having less than 2 neighboring candidate pixels. This results in a trimming back of candidate pixel regions 53-57, as shown.

Component-connecting block 15 of Fig. 1, then formalizes the cohesive groups of candidate pixels making up regions 53-57 into respective components. As explained above, at this point, statistical data regarding each respective component is generated and stored for later use. In the present example, the results of such calculation would indicate that the pixel having brightest value would lie within one of regions 53 and 55. In the present example, it is assumed that region, or component, 53 includes the pixel with the highest bright parameter, and thus first redeye identifier block 17 of Fig. 1 would identify component 53 as the first redeye region.

At this point, second identifier block 19 surveys the statistical data of remaining components 55 and 57. Although the statistical data would show that both components 55 and 57 have similar average bright parameters, region 57 is disqualified do to it having a size outside a range defined by the size of component 53. Component 55, whoever, does have a size within a range defined by component 53 and thus it is identified as a candidate component for possibly being a redeye region. In the present example, there are no other candidate components and region 55 is thus identified as the second pupil. If additional candidate components had been identified, however, then their respective boundary square population

value would be compared and the most densely populated candidate component would be targeted as being the second redeye region.

The resultant matrix 70 following second redeye identifier block 19 is shown in Fig. 5. As shown, all components except for the identified redeye regions 53 and 55 are removed.

With reference to Fig. 6, a second matrix 50' generated either following block 19 or generated concurrently with matrix 50 by matrix generator block 11, compares the pixels of the selection rectangle 37 of Fig. 2 with a relaxed threshold based on the threshold used to generate matrix 50. Since the threshold used to generate matrix 50' of Fig. 6 is lower, it shows the inclusion of enlarged regions 53' and 55', and also many more candidate pixels 51' correspond to reddish tones in glass 41 and candidate pixels 57' correspond to bright red cloth 43.

As explained above, the area of candidate components 53 and 55 of matrix 70 in Fig. 5 are dilated outward and are bit-wise ANDed with corresponding candidate pixels 53' and 57' of matrix 50'. This operation is explained above as being conducted by final redeye matrix generator block 25 of Fig. 1.

With reference to Fig. 7, the resultant matrix 80 following block 25 is then use for redeye correction. Pupil regions 83 and 87 define the result of the above mentioned AND operation, and are shown dark to indicate that these regions of the original selection rectangle image are replaced with a uniform dark gray color. The pixels adjacent regions 83 and 87 constitute respective feathering region 83' and 87', and are shown light gray to indicate that these regions of the original selection rectangle image are replaced with a uniform gray color of higher brightness than that of respective regions 83 and 87.

With reference to Fig. 8, the corrected image shows that pupil regions 33 and 35 have been replaced with a combination of dark and gray pixels while maintaining clearly visible glint regions 39 and 38.

While the invention has been described in conjunction with several specific embodiments, it is evident to those skilled in the art that many further alternatives, modifications and variations will be apparent in light of the foregoing description. Thus, the invention described herein is intended to embrace all such alternatives, modifications, applications and variations as may fall within the spirit and scope of the appended claims.